

IN THE CLAIMS

What is claimed is:

5

- 1. An apparatus, comprising:
 - at least one input port capable of being coupled to at least one substantially simultaneously preselected first other computation node, the input port to receive input data;
 - 10 a first store coupled to the at least one input port to store the input data;
 - a second store coupled to an external instruction sequencer, the second store to receive and store an instruction from the external instruction sequencer;
 - an instruction wakeup unit to match the input data to the instruction;
 - 15 at least one execution unit to execute the instruction using the input data to produce output data; and
 - at least one output port capable of being coupled to at least one substantially simultaneously preselected second other computation node.

20 2. The apparatus of claim 1, further comprising:

- a router to direct the output data from the at least one output port to the at least one substantially simultaneously preselected second other computation node.

25 3. The apparatus of claim 2, wherein the instruction includes a destination address associated with the at least one substantially simultaneously preselected second other computation node, and wherein the router is capable of using the destination address to direct the output data to the at least one substantially simultaneously preselected second other computation node.

30

4. The apparatus of claim 3, wherein the destination address is generated by a mechanism selected from the group consisting of: a compiler and a run-time trace mapper.
5. The apparatus of claim 2, wherein the instruction includes a destination address associated with the computation node, and wherein the router is capable of using the destination address to direct the output data to the computation node.
- 10 6. The apparatus of claim 1, wherein the execution unit comprises at least one calculation module selected from the group consisting of: an arithmetic logic unit, a floating point unit, a memory address unit, and a branch unit.
- 15 7. The apparatus of claim 1, wherein the second store is capable of storing multiple instructions.
8. The apparatus of claim 1, wherein the first store is capable of storing multiple operands.
- 20 9. The apparatus of claim 1, wherein the at least one output port is coupled to a direct channel, and wherein an input port of the at least one substantially simultaneously preselected second other computation node is coupled to the direct channel.
- 25 10. The apparatus of claim 1, wherein the at least one input port is coupled to a direct channel, and wherein an output port of the at least one substantially simultaneously preselected first other computation node is coupled to the direct channel.

11. A system, comprising:
 - an external instruction sequencer to fetch a group of instructions including an instruction; and
 - a first preselected computation node including at least one input port
 - 5 capable of being coupled to at least one first other preselected computation node,
 - the input port to receive input data, a first store coupled to the at least one input port to store the input data, a second store coupled to the instruction sequencer, the second store to receive and store the instruction, an instruction wakeup unit to match the input data to the instruction, at least one execution unit to execute
 - 10 the instruction using the input data to produce output data, at least one output port capable of being coupled to at least one second other preselected computation node, and a router to direct the output data from the at least one output port to the at least one second other preselected computation node.
- 15 12. The system of claim 11, further comprising:
 - a second preselected computation node including at least one input port capable of being coupled to at least one third other preselected computation node, the input port to receive other input data, a first store coupled to the at least one input port to store the other input data, a second store coupled to the
 - 20 instruction sequencer, the second store to receive and store an other instruction selected from the group of instructions, an instruction wakeup unit to match the other input data to the other instruction, at least one execution unit to execute the other instruction using the other input data to produce other output data, at least one output port capable of being coupled to at least one fourth other preselected computation node, and a router to direct the other output data from the at least one output port to the at least one fourth other preselected computation node.
- 25 13. The system of claim 12, further comprising:
 - a register file to receive indications to send operands to be used by
 - 30 instructions at the first and the second preselected computations nodes.

14. The system of claim 12, wherein the output port of the first preselected computation node is coupled to the input port of the second preselected computation node.

5 15. The system of claim 11, further comprising:
a grid of computation nodes including the first preselected computation node, wherein the grid of computation nodes includes M rows of computation nodes, and wherein each one of the M rows of computation nodes includes an instruction cache.

10 16. The system of claim 11, further comprising:
an instruction memory coupled to the instruction sequencer, the instruction memory to store the group of instructions.

15 17. The system of claim 11, further comprising:
a block termination control module to detect execution termination of the group of instructions; and
a register file coupled to the block termination control module.

20 18. A method, comprising:
partitioning a program into a plurality of groups of instructions;
assigning a group of instructions selected from the plurality of groups of instructions to a plurality of interconnected preselected computation nodes;
loading the group of instructions to the plurality of interconnected preselected computation nodes; and
executing the group of instructions as each one of the instructions in the group of instructions receives all necessary associated operands for execution.

19. The method of claim 18, wherein at least one computation node included in
30 the plurality of interconnected preselected computation nodes has at least one input port capable of being coupled to at least one preselected first other

computation node included in the plurality of interconnected preselected computation nodes, the input port to receive input data, a first store coupled to the at least one input port to store the input data, a second store coupled to an instruction sequencer, the second store to receive and store the at least one instruction, an instruction wakeup unit to match the input data to the at least one instruction, at least one execution unit to execute the at least one instruction using the input data to produce output data, at least one output port capable of being coupled to at least one second other preselected computation node included in the plurality of interconnected preselected computation nodes, and a router to direct the output data from the at least one output port to the at least one preselected second other computation node.

20. The method of claim 18, wherein at least one of the plurality of groups of instructions is a basic block.

15 21. The method of claim 18, wherein at least one of the plurality of groups of instructions is a hyperblock.

22. The method of claim 18, wherein at least one of the plurality of groups of 20 instructions is a superblock.

23. The method of claim 18, wherein at least one of the plurality of groups of instructions is an instruction trace constructed by a hardware trace construction unit at run time.

25 24. The method of claim 18, wherein loading the group of instructions to the plurality of interconnected preselected computation nodes includes:

30 sending at least two instructions selected from the group of instructions from an instruction sequencer to a selected computation node included in the plurality of interconnected preselected computation nodes for storage in a store.

25. The method of claim 18, wherein executing the group of instructions as each one of the instructions in the group of instructions receives all necessary associated operands for execution includes:

matching at least one instruction selected from the group of instructions

5 with at least one operand received from an other computation node included in

-- the plurality of interconnected preselected computation nodes.

26. The method of claim 18, wherein loading the group of instructions to the plurality of interconnected preselected computation nodes includes:

10 sending a first set of instructions selected from a first group of instructions selected from the plurality of groups of instructions from an instruction sequencer to the plurality of interconnected preselected computation nodes for storage in a first frame included in a first computation node included in the plurality of interconnected preselected computation nodes; and

15 sending a second set of instructions selected from the first group of instructions from the instruction sequencer to the plurality of interconnected preselected computation nodes for storage in a second frame included in the first computation node.

20 27. The method of claim 18, wherein assigning a group of instructions selected from the plurality of groups of instructions to a plurality of interconnected preselected computation nodes includes:

assigning a first group of instructions to a first set of frames included in the plurality of interconnected preselected computation nodes;

25 assigning a second group of instructions to a second set of frames included in the plurality of interconnected preselected computation nodes, wherein the first group and the second group of instructions are capable of concurrent execution, and wherein at least one output datum associated with the first group of

30 instructions is written to a register file and passed directly to the

- second group of instructions for use as an input datum by the second group of instructions.

28. An article comprising a machine-accessible medium having associated data,
5 wherein the data, when accessed, results in a machine performing:

- partitioning a program into a plurality of groups of instructions;
- assigning a group of instructions selected from the plurality of groups of instructions to a plurality of interconnected preselected computation nodes;
- loading the group of instructions to the plurality of interconnected 10 preselected computation nodes; and
- executing the group of instructions as each one of the instructions in the group of instructions receives all necessary associated operands for execution.

29. The article of claim 28, wherein partitioning the program into the plurality of 15 groups of instructions is performed by a compiler.

30. The article of claim 28, wherein partitioning the program into the plurality of groups of instructions is performed by a run-time trace mapper.

20 31. The article of claim 28, wherein the machine-accessible medium further includes data, which when accessed by the machine, results in the machine performing:

- statically assigning all of the plurality of groups of instructions for execution.

25 32. The article of claim 31, wherein the machine-accessible medium further includes data, which when accessed by the machine, results in the machine performing:

- dynamically issuing one or more instructions from at least one of the 30 plurality of groups of instructions for execution.

33: The article of claim 28, wherein the machine-accessible medium further includes data, which when accessed by the machine, results in the machine performing:

generating a wakeup token to reserve an output data channel to connect

5 selected computation nodes included in the plurality of interconnected

-- preselected computation nodes.

34. The article of claim 28, wherein the machine-accessible medium further includes data, which when accessed by the machine, results in the machine

10 performing:

detecting execution termination of the group of instructions including an

output having architecturally visible data; and

committing the architecturally visible data to a register file.

15 35. The article of claim 28, wherein the machine-accessible medium further includes data, which when accessed by the machine, results in the machine

performing:

detecting execution termination of the group of instructions including an

output having architecturally visible data; and

20 committing the architecturally visible data to a memory.

36. The article of claim 28, wherein the machine-accessible medium further includes data, which when accessed by the machine, results in the machine

performing:

25 routing an output datum arising from executing the group of instructions

to a consumer node included in the plurality of interconnected preselected

computation nodes, wherein the address of the consumer node is included in a

token associated with at least one instruction included in the group of

instructions.